


REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	8/8/95	6/20-91-12/31/94	
4. TITLE AND SUBTITLE		5. FUNDING NUMBERS	
Continuation of the POSTGRES Project		DAAL03-91-0183	
6. AUTHOR(S)		7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	
Dr. Michael Stonebraker		<p>University of California Electronics Research Laboratory CS Division 617 Soda Hall Berkeley, CA 94720-1776</p> 	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		8. PERFORMING ORGANIZATION REPORT NUMBER	
<p>U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211</p>		10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES			
<p>The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.</p>			
12a. DISTRIBUTION / AVAILABILITY STATEMENT		12b. DISTRIBUTION CODE	
Approved for public release; distribution unlimited.			
13. ABSTRACT (Maximum 200 words)			
<p>During this grant we have focused on four tasks related to the POSTGRES project, namely: 1) refinement and implementation of the POSTGRES rules system; 2) integration of tertiary memory support into POSTGRES; 3) efficient support for very large arrays; 4) efficient support for expensive predicates. In this report, we discuss each in turn.</p>			
19951005 050			
DTIC QUALITY INSPECTED 5			
14. SUBJECT TERMS			15. NUMBER OF PAGES
databases, DBMS, relational databases			4
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

CONTINUATION OF THE POSTGRES PROJECT

FINAL REPORT

MICHAEL STONEBRAKER

AUGUST 8, 1994

U.S. ARMY RESEARCH OFFICE

DAALD3-91-G-0183

UNIVERSITY OF CALIFORNIA
BERKELEY, CA 94720

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

Continuation of the POSTGRES Project

Final Report

*Michael Stonebraker, Lawrence Rowe
EECS Dept.
University of California, Berkeley*

1. Introduction

During this grant we have focused on four tasks related to the POSTGRES project, namely:

- 1) refinement and implementation of the POSTGRES rules system
- 2) integration of tertiary memory support into POSTGRES
- 3) efficient support for very large arrays
- 4) efficient support for expensive predicates

In the rest of this report, we discuss each topic in turn.

2. Rules System

Near the end of the previous ARO grant, we developed a collection of algorithms for supporting virtual classes in POSTGRES as well as alternate versions of classes. This work was published in [STON90]. Briefly, we discovered that both functions could be implemented by innovative uses of our POSTGRES rules system. As such, the special purpose, low-level, code required in traditional systems to support these constructs can be replaced by a small collection of rules written in the POSTGRES rule language.

To validate the utility of this idea, we have implemented the rules for both systems in POSTGRES and examined their performance. Specifically, we have found that performing version management using the POSTGRES rules approach is competitive with utilizing previous low-level code we had written in the mid 1980's. Moreover, it is much easier to implement and modify a rules-based system than one based on hard code in a 3rd generation programming language.

In addition, our current implementations for rules provide "immediate activation", i.e. the action for each rule is triggered at the time the event specified in the rule becomes true. At times, a user would like "deferred execution", i.e. he would like rule activation to be delayed until the commit time of an enclosing transaction. We have investigated how to perform deferred execution without having to maintain complex bookkeeping about the effects of a transaction during execution. Unfortunately, there does not seem to be an easy way to implement this functionality without extensive reworking of the current code base. A paper on the options and problems in this area along with some suggestions for future investigation appeared in the *IEEE Transactions on Knowledge and Data Engineering* [STON92].

Lastly, we have worked on applying DBMS rules systems to a substantial real world problem to validate the concepts. In particular, we have implemented the notion of calendars for a financial services time series application using the POSTGRES rules system. A report on this matter appeared in the 1994 IEEE Data Engineering conference [CHAN94].

3. Tertiary Memory Support

We have worked on two separate problems in this area, namely the efficient integration of tertiary memory support in a DBMS query optimizer and the provision of a file system interface on top of a DBMS. We discuss each topic in turn.

First, we have performed a detailed study of query optimization in a tertiary memory context. Specifically, we have examined query processing algorithms implemented by DBMSs to perform restrictions, projections and joins and have discovered versions of these algorithms optimized for tertiary memory. Moreover, we have investigated the optimal scheduling of the robot arm between tertiary memory data and a CPU. We have found that substantial performance improvements are available through careful scheduling of the "batch" of requests that are outstanding in a multi-user environment. A report on this matter has been accepted at the 1995 Very Large Data Base Conference [SARA95].

Second, we have proposed that a standard operating system file system interface be simulated on top of a DBMS-managed storage hierarchy. In current systems, the DBMS must exist on top of the file system, and serious function and performance consequences result. This has led many commercial DBMSs to bypass the file system and implement their systems directly on top of a "raw" disk. In contrast, it is possible to reverse the two systems, and implement a file system on top of a DBMS, a concept which we called "Inversion".

Using the Inversion concept, any file system operation would turn into a query to the DBMS. If a user performs very small reads and writes, then the performance of this approach may be problematic. However, if a user is reading or writing very large objects, often to tertiary memory, then there should be little, if any, performance difference between the two approaches.

We have implemented a prototype version of this Inversion concept, and a report on this topic appeared in the 1993 IEEE Data Engineering Conference [STON93]. It demonstrates that very reasonable performance is available as long as reads and writes involve large objects.

4. Storage of Multidimensional Arrays

Next, we have investigated the layout of very large multidimensional arrays on secondary and tertiary memory. Specifically, in the companion Sequoia 2000 project, we have used POSTGRES to support the DBMS needs of a collection of atmospheric science users of General Circulation Models (GCMs). They wish to store the output of their models, which is in the form of a four dimension array, in a data base. Moreover, subsequently, they wish to form various projections of this array data, an operation often called "creating a hyperslab".

We have discovered for typical hyperslab workloads that storing arrays in "Fortran order" is very inefficient. Furthermore, "chunking" the array shows a marked speedup. A paper on this topic appeared in the 1994 IEEE Data Engineering Conference [SARA94].

5. Optimization of Expensive Functions

We have extended the POSTGRES optimizer to deal with functions which are expensive to compute. For example, consider the following query:

```
retrieve (EMP.name)
  where beard (EMP.picture) = "red" and EMP.age < 30
```

In this case, the first clause in the predicate consumes perhaps 100 million CPU instructions to perform a pattern analysis of the image to determine if the picture is of a person with a beard. Moreover, the function must read a megabyte or more of data in the process. In contrast, the second clause requires perhaps 100 instructions and the reading of four bytes. When there are dramatic differences between the computational demands of the various clauses in the predicate, it is crucial for an optimizer to be "smart" about the processing order of the clauses. In addition, the optimizer should consider delaying the processing of clauses involving expensive functions as long as possible when constructing the query plan.

We demonstrated sketchy results in [STON91] of an approach to this problem. A more extensive analysis of the topic appeared in the 1993 ACM-SIGMOD annual conference [HELL93]. At the current

For	
<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>
tion	
Low	
Availability Codes	
Dist	Avail and/or Special
A	

time, these algorithms have been fully implemented in the POSTGRES DBMS, as well as in the commercial version of the code line, a product called Illustra.

6. Papers Published Under This Grant

[CHAN94] Chandra, R. et. al., "Implementing Calendars and Temporal Rules in Next Generation Databases," Proc. 1994 IEEE Data Engineering Conference, Houston, Tx., Feb. 1994.

[HELL93] Hellerstein, J. and Stonebraker, M., "Predicate Pushdown for Expensive Functions," Proc. 1993 ACM-SIGMOD Conference on Management of Data, Philadelphia, Pa., May 1993.

[SARA94] Sarawagi, S. and Stonebraker, M., "Efficient Organization of Large Multidimensional Arrays," Proc. 1994 IEEE Data Engineering Conference, Houston, Tx., Feb. 1994.

[SARA95] Sarawagi, S., "Query Processing in Tertiary Memory Databases," Proc. 1995 VLDB Conference, Zurich, Switzerland, Sept. 1995.

[STON90] Stonebraker, M. et. al., "On Rules, Procedures, Caching, and Views," Proceedings of the ACM-SIGMOD Conference on Management of Data, Atlantic City, N.J., June 1990.

[STON91] Stonebraker, M., "Managing Persistent Objects in a Multi-level Store," Proceedings of the ACM-SIGMOD Conference on Management of Data, Denver, Co., June 1991.

[STON92] Stonebraker, M., "The Integration of Rule Systems and Data Base Systems," IEEE Transaction on Knowledge and Data Engineering, October 1992.

[STON93] Stonebraker, M. and Olson, M., "Large Object Support in POSTGRES," Proc. 1993 IEEE Data Engineering Conference, Vienna, Austria, April 1993.

7. Scientific Personnel

Robert Devine -- MS

Joseph Hellerstein -- MS

Anant Jhingram -- Ph.D.

Spyros Potamianos -- Ph.D.

8. Report of Inventions

See Attached.

REPORT OF INVENTIONS AND SUBCONTRACTS

(Pursuant to "Patent Rights" Contract Clause) (See Instructions on Reverse Side.)

				<p style="text-align: center;">Form Approved OMB No. 0704-0297 Expires Jun 30, 1992</p> <p>Public reporting burden for this collection of information is estimated to average 5 minutes per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1220 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0297), Washington, DC 20503.</p>			
1a. NAME OF CONTRACTOR / SUBCONTRACTOR University of California Electronics Research Lab 617 Soda Hall Berkeley, CA 94720-1776		1b. CONTRACT NUMBER DAAL03-91-C-0183 1c. ADDRESS (Include ZIP Code) P.O. Box 12211 Research Triangle Park, NC		2a. NAME OF GOVERNMENT PRIME CONTRACTOR U.S. Army Research Office 2b. ADDRESS (Include ZIP Code) 620/91-12/31/94 2c. CONTRACT NUMBER DAAL03-91-C-0183 2d. AWARD DATE (YYMMDD) 6/20/91-12/31/94			
3. TYPE OF REPORT (X one)		<input type="checkbox"/> a. INTERIM <input checked="" type="checkbox"/> b. FINAL					
4. REPORTING PERIOD (YYMMDD)		5. FROM 6/20/91 TO 12/31/94					
SECTION I - SUBJECT INVENTIONS							
6. SUBJECT INVENTIONS REQUIRED TO BE REPORTED BY CONTRACTOR / SUBCONTRACTOR (If "None," so state)							
6a. NAME(S) OF INVENTOR(S) (Last, First, MI)		6b. TITLE OF INVENTION(S)		6c. DISCLOSURE NO., PATENT APPLICATION SERIAL NO. OR PATENT NO. 6d. ELECTION TO FILE PATENT APPLICATIONS 6e. CONFIRMATORY INSTRUMENT OR ASSIGNMENT FORWARDED TO CONTRACTING OFFICER			
6e. CONFIRMATORY INSTRUMENT OR ASSIGNMENT FORWARDED TO CONTRACTING OFFICER		<input type="checkbox"/> (1) United States <input type="checkbox"/> (2) Foreign <input type="checkbox"/> (a) Yes <input type="checkbox"/> (b) No <input type="checkbox"/> (a) Yes <input type="checkbox"/> (b) No <input type="checkbox"/> (1) Yes <input type="checkbox"/> (2) No					
6f. EMPLOYER OF INVENTOR(S) NOT EMPLOYED BY CONTRACTOR / SUBCONTRACTOR		6g. ELECTED FOREIGN COUNTRIES IN WHICH A PATENT APPLICATION WILL BE FILED					
6g. ELECTED FOREIGN COUNTRIES IN WHICH A PATENT APPLICATION WILL BE FILED		6h. FOREIGN COUNTRIES OF PATENT APPLICATION					
6i. SUBCONTRACTS AWARDED BY CONTRACTOR / SUBCONTRACTOR (If "None," so state)		6j. SUBCONTRACT COUNTRIES OF PATENT APPLICATION					
6k. SUBCONTRACTS AWARDED BY CONTRACTOR / SUBCONTRACTOR (If "None," so state)		6l. SUBCONTRACT COUNTRIES OF PATENT APPLICATION					
7. CERTIFICATION OF REPORT BY CONTRACTOR / SUBCONTRACTOR		7a. NAME OF AUTHORIZED CONTRACTOR / SUBCONTRACTOR OFFICIAL (Last, First, MI) Stonebraker, Michael 7b. TITLE Principal Investigator					
7c. I certify that the reporting party has procedures for prompt identification and timely disclosure of "Subject Inventions," that such procedures have been followed and that all "Subject Inventions" have been reported.		7d. SUBCONTRACT DATES (YYMMDD)					
7e. DATE SIGNED		7f. SUBCONTRACT DATES (YYMMDD)					
7g. DATE SIGNED		7h. SUBCONTRACT DATES (YYMMDD)					
SECTION II - SUBCONTRACTS (Containing a "Patent Rights" clause)							
SECTION III - CERTIFICATION							
8. NAME OF SUBCONTRACTOR(S)		8a. ADDRESS (Include ZIP Code) 8b. SUBCONTRACT NO.(S) 8c. DATE (YYMM) 8d. DESCRIPTION OF WORK TO BE PERFORMED UNDER SUBCONTRACT(S)					
8e. DESCRIPTION OF WORK TO BE PERFORMED UNDER SUBCONTRACT(S)		8f. SUBCONTRACT DATES (YYMMDD)					
8g. SUBCONTRACT DATES (YYMMDD)		8h. SUBCONTRACT DATES (YYMMDD)					
Previous editions are obsolete.							

DD FORM 882 INSTRUCTIONS

GENERAL

This form is for use in submitting INTERIM and FINAL invention reports to the Contracting Officer and for use in the prompt notification of the award of subcontracts containing a "Patent Rights" clause. If the form does not afford sufficient space, multiple forms may be used or plain sheets of paper with proper identification of information by Item Number may be attached.

An INTERIM report is due at least every 12 months from the date of contract award and shall include (a) a listing of "Subject Inventions" during the reporting period, (b) a certification of compliance with required invention identification and disclosure procedures together with a certification of reporting of all "Subject Inventions," and (c) any required information not previously reported on subcontracts awarded during the reporting period and containing a "Patent Rights" clause.

A FINAL report is due within 6 months if contractor is a small business firm or domestic nonprofit organization and within 3 months for all others after completion of the contract work and shall include (a) a listing of all "Subject Inventions" required by the contract to be reported, and (b) any required information not previously reported on subcontracts awarded during the course of or under the contract and containing a "Patent Rights" clause.

While the form may be used for simultaneously reporting inventions and subcontracts, it may also be used for reporting, promptly after award, subcontracts containing a "Patent Rights" clause.

Dates shall be entered where indicated in certain items on this form and shall be entered in four or six digit numbers in the order of year and month (YYMM) or year, month and day (YYMMDD). Example: April 1986 should be entered as 8604 and April 15, 1986 should be entered as 860415.

Item 1a. Self-explanatory.

Item 1b. Self-explanatory.

Item 1c. If "same" as item 2c, so state.

Item 1d. Self-explanatory.

Item 2a. If "same" as item 1a, so state.

Item 2b. Self-explanatory.

Item 2c. Procurement Instrument Identification (PII) number of contract (DFAR 4.7003).

Item 2d thru 5e. Self-explanatory.

Item 5f. The name and address of the employer of each inventor not employed by the contractor or subcontractor is needed because the Government's rights in a reported invention may not be determined solely by the terms of the "Patent Rights" clause in the contract.

Example 1: If an invention is made by a Government employee assigned to work with a contractor, the Government rights in such an invention will be determined under Executive Order 10096.

Example 2: If an invention is made under a contract by joint inventors and one of the inventors is a Government employee, the Government's rights in such an inventor's interest in the invention will also be determined under Executive Order 10096, except where the contractor is a small business or nonprofit organization, in which case the provisions of Section 202(e) of P.L. 96-517 will apply.

Item 5g (1). Self-explanatory.

Item 5g (2). Self-explanatory with the exception that the contractor or subcontractor shall indicate, if known at the time of this report, whether applications will be filed under either the Patent Cooperation Treaty (PCT) or the European Patent Convention (EPC). If such is known, the letters PCT or EPC shall be entered after each listed country.

Item 6a. Self-explanatory.

Item 6b. Self-explanatory.

Item 6c. Self-explanatory.

Item 6d. Patents Rights Clauses are located in FAR 52.227.

Item 6e thru 7b. Self-explanatory.

Item 7c. Certification not required by small business firms and domestic nonprofit organizations.